Market Practitioners Group

The Jibar transition: **The last mile**

Event-Aware Jump-Diffusion for the JIBAR–ZARONIA Spread

Mesias Alfeus[a,]

[a]*Department of Statistics and Actuarial Science,*
University of Stellenbosch, South Africa.
§National Institute for Theoretical & Computational Sciences (NITheCS)

Abstract

We study the JIBAR–ZARONIA basis with a square-root (CIR) diffusion augmented by both scheduled monetary policy jumps and data-driven jump detections. The design is *event-aware* in two senses. First, intraday Monetary Policy Committee (MPC) announcements are aligned to end-of-day observations using a day-tolerance τ and $\pm K$-bar windows to reconcile announcement timing, holidays, and market close. Second, residual discontinuities are identified via threshold rules that distinguish *step-like* shifts from *spike-and-reversal* episodes (Melanson and Longtin, 2019). Parameters are estimated by feasible generalized least squares (GLS), which respects the CIR heteroskedasticity $\mathrm{Var}(\Delta X_t \mid X_{t-1}) \propto X_{t-1}\Delta t$, and the fitted dynamics are translated into risk through simulation-based Value-at-Risk (VaR) and Potential Future Exposure (PFE). Empirically, the spread exhibits strong mean reversion toward a lower steady state, small average scheduled jump effects at policy dates once mean reversion is controlled, and non-negligible tail risk under clustered shocks. Overlap metrics show that absolute-exceedance detection aligns best with policy *steps*, whereas sign-flip excels for transient *spikes*. The framework is transparent, replicable, and directly applicable to fallback design, discounting, and risk governance during South Africa's benchmark transition, offering a practical bridge between event calendars, econometric identification, and risk measurement.

Keywords: Benchmark transition, JIBAR, ZARONIA, CIR model, Jump detection, VaR, Potential Future Exposure, South Africa

Email address: mesias@sun.ac.za (Mesias Alfeus)

1. Introduction

Post-crisis benchmark reform replaces panel-based IBORs with transaction-based overnight rates to improve representativeness, governance, and robustness (International Organization of Securities Commissions (IOSCO), 2021; Alternative Reference Rates Committee (ARRC), 2023; Federal Reserve Bank of New York, 2021). In South Africa, the forward-looking unsecured term rate JIBAR is being retired in favor of ZARONIA, a transaction-based overnight index grounded in observed money-market activity (South African Reserve Bank (SARB), 2023). The transition affects valuation (discounting curves), hedge effectiveness, collateral and margin, and the treatment of legacy exposures through fallback spreads (Bloomberg, 2021). Beyond policy and operations, the key quantitative object is the JIBAR–ZARONIA spread: it embeds credit/liquidity premia and microstructure features that matter for pricing and risk management.

Two design facts motivate our modeling. First, policy information arrives on a known calendar (MPC announcements), but the data we use are end-of-day—so any econometric treatment must reconcile different clocks. Second, while a diffusion captures persistent volatility, the spread also experiences discontinuities (policy steps, liquidity shocks). These observations suggest a CIR jump–diffusion with explicit event alignment and a complementary, data-driven detector of unscheduled jumps.

We offer a compact, production-ready framework that (i) couples a CIR diffusion with scheduled policy jumps (Backwell and Hayes, 2022) mapped to the daily grid via a tolerance τ and $\pm K$ window; (ii) adds a threshold-based detector that distinguishes step-like moves from spike-and-reversal patterns (Melanson and Longtin, 2019); (iii) estimates parameters by feasible GLS, consistent with the CIR variance $\mathrm{Var}(\Delta X_t \mid X_{t-1}) \propto X_{t-1}\Delta t$; and (iv) delivers VaR and PFE on the empirical calendar. A simple set of overlap metrics (precision, recall, F1, Jaccard) then quantifies how scheduled and detected jumps co-locate under (τ, K) choices.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature. Section 3 describes the data and variable construction. Section 4 sets out the methodological framework. Section 5 presents the empirical results and robustness checks. Section 6 concludes and discusses implications.

2. Related Literature

A large literature documents the weaknesses of IBOR benchmarks and the post-crisis migration to transaction-based overnight risk-free rates (RFRs). The LIBOR episode exposed both embedded bank credit premia and governance vulnerabilities

in panel submissions (Backwell et al., 2025; Ashton and Christophers, 2015; Kiff, 2012). In response, IOSCO articulated principles for robust benchmarks emphasising data quality, transparency, and governance (International Organization of Securities Commissions (IOSCO), 2021), and policy groups developed transition toolkits and fallbacks, notably the ARRC's corpus for SOFR adoption (Alternative Reference Rates Committee (ARRC), 2023) and the Bloomberg rule book implementing contractual fallback mechanics (Bloomberg, 2021). The Federal Reserve Bank of New York (FRBNY's) guidance consolidates market conventions for SOFR compounding and use in discounting (Federal Reserve Bank of New York, 2021).

Within this global shift, research evaluates whether SOFR improves upon LIBOR for pricing and risk management, with evidence favouring transaction-based, nearly risk-free rates for discounting and collateral frameworks (Jermann, 2019; Klingler and Syrstad, 2024). Risk measurement under benchmark reform continues to rely on established tools such as Value-at-Risk (VaR) and Expected Shortfall (ES), with methodological references including Jorion (1996), and Best (2000). These frameworks remain applicable but require attention to basis risk during coexistence of legacy and new rates.

South Africa's pathway mirrors global reforms but features local institutions and timelines. JIBAR, a forward-looking unsecured term rate constructed from bank submissions, raised questions about representativeness and resilience even absent proven manipulation (de Jager et al., 2013). ZARONIA, by contrast, is a transaction-based overnight rate designed and phased in by the SARB with detailed market conventions and implementation guidance (South African Reserve Bank (SARB), 2023). The transition plan outlines observation, adoption, and active migration phases leading to the cessation of JIBAR (Jibar Transition Plan Team, 2024). Commentary and industry surveys highlight operational readiness, systems reconfiguration, liquidity development, and communication challenges, alongside broad support for the transparency benefits of ONIAs (MPG Workstream, 2024; Nyandeni, 2024; Monocle Solutions, 2023; Digiata Team, 2024; Santos, 2024; Ilkova and Silberman, 2024).

A focal object for valuation and fallback design is the spread between legacy IBORs and new ONIAs. International fallbacks rely on compounded overnight rates plus historical spread adjustments to approximate term-credit components, with design choices (e.g., median versus mean windows) trading robustness against responsiveness (Bloomberg, 2021; Alternative Reference Rates Committee (ARRC), 2023). In the South African context, SARB proposes a JIBAR fallback methodology aligned with these principles while acknowledging data-history and market-structure specifics. Recent applied work synthesises progress and open questions for the JIBAR–ZARONIA

transition, including impacts on valuation, hedging, and governance (Alfeus, 2024).

Methodologically, modelling short-rate and spread dynamics with state-dependent volatility often leverages square-root diffusions augmented with discontinuities. Data-driven jump detection complements parametric jump-diffusion modelling; for stationary processes, threshold-based identification with tail-balanced selection and simple jump-time rules can recover spike-and-reversal and step-change episodes (Melanson and Longtin, 2019). In the transition setting, combining scheduled policy events (e.g., MPC announcements) with data-driven detections provides a practical, event-aware approach to quantify policy-day effects, unscheduled shocks, and their risk implications for VaR/PFE.

3. Data and Measures

We gathered daily 3M JIBAR, ZARONIA (proxy from 2016 - 2022 provided by SARB), and compute the spread. JIBAR Sample: 04/01/2016 to 31/08/2024 is obtained from Bloomberg. We preserve calendar gaps via $\Delta t_i = (t_i - t_{i-1})/365$, remove non-finite values, and align time stamps.

MPC announcement dates are obtained from SARB calendars. We map each date to the nearest trading day within $\tau \in \{2, 3, 4\}$ days.

The spread has mean 44.15 bp and sd 17.12 bp; skewness -0.2808; kurtosis 2.255. Jarque–Bera rejects normality at 1%.

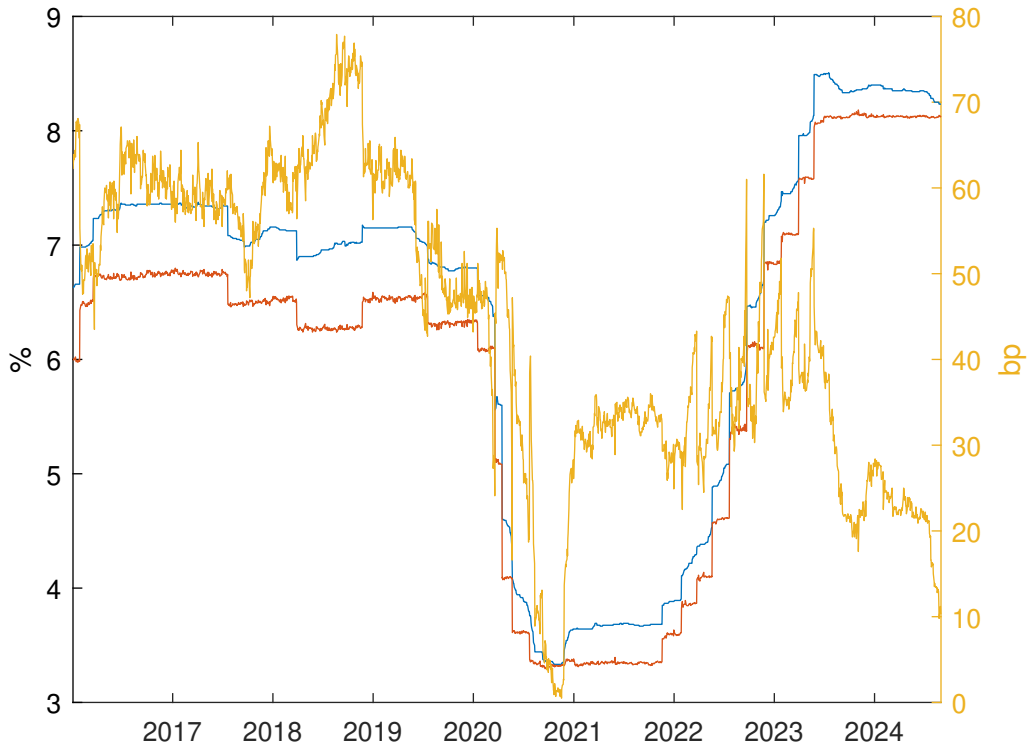

Figure 1: JIBAR, ZARONIA, and the JIBAR–ZARONIA spread time series

Table 1: Descriptive statistics for JIBAR, ZARONIA, and their spread.

Series	Mean (bp)	Std (bp)	Skew.	Kurt.	JB	$\rho\big(\lvert\Delta X\rvert, \sqrt{X}\big)$
JIBAR	641.1	162.1	−0.750	2.158	267.00***	−0.028**
ZARONIA	597.0	156.5	−0.548	2.141	175.00***	−0.006**
Spread	44.15	17.12	−0.281	2.256	78.45***	0.088***

Notes: Rates are in basis points (bp). "Spread" is JIBAR − ZARONIA. JB = Jarque–Bera statistic computed from sample skewness and kurtosis; large values imply non-normal increments (here, $p < 1\%$ in all cases). $\rho\big(\lvert\Delta X\rvert, \sqrt{X}\big)$ is the Pearson correlation between absolute one-step changes and the square root of the level, a simple proxy for level-dependent volatility. *, **, *** denotes the 10%, 5% and 1% level of significance, respectively.

4. Event-Aware Model and Estimation

4.1. Set-up and Notation

Let X_t denote the JIBAR–ZARONIA spread observed on a (possibly irregular) daily grid $\{t_i\}_{i=0}^{N}$ with $t_i \in \mathbb{R}$ (calendar time). Define one-step increments $\Delta X_{t_i} := X_{t_i} - X_{t_{i-1}}$ and calendar step lengths $\Delta t_i := (t_i - t_{i-1})/365$. Monetary Policy Committee (MPC) announcement dates are provided as a calendar set $\mathcal{S} := \{s_k\}_{k=1}^{K}$.

We model $X_t \geq 0$ as a mean-reverting square-root diffusion augmented by discrete jumps occurring on MPC dates mapped to the observation grid. This "event-aware" design lets us quantify both continuous dynamics and policy-day displacements.

4.2. Continuous Dynamics with Scheduled Jumps

The continuous-time specification is

$$\mathrm{d}X_t = \kappa(\theta - X_t)\,\mathrm{d}t + \sigma\sqrt{X_t}\,\mathrm{d}W_t + J_t\,\mathrm{d}N_t, \tag{1}$$

where $\kappa > 0$ is the speed of mean reversion, $\theta > 0$ the long-run level, and $\sigma > 0$ the diffusion scale. The counting process N_t is *scheduled*: it jumps when t coincides with an MPC announcement mapped to the sample grid (Section 4.3). The jump size J_t is allowed to be signed and time-varying, but for estimation parsimony we summarise its one-step mean effect by a constant μ_J.

On $\{t_i\}$, a quasi-Gaussian Euler discretisation of (1) yields

$$\Delta X_{t_i} = \kappa(\theta - X_{t_{i-1}})\,\Delta t_i + \mu_J I_i + \sigma\sqrt{X_{t_{i-1}}\,\Delta t_i}\,\varepsilon_i, \qquad \varepsilon_i \sim \mathcal{N}(0,1), \tag{2}$$

where $I_i \in \{0,1\}$ indicates whether t_i is mapped to an MPC event. The conditional variance of increments is proportional to $X_{t_{i-1}}\Delta t_i$, producing state-dependent heteroskedasticity.

4.3. Mapping MPC Dates to the Observation Grid

Announcements are intraday; the data are end-of-day. We therefore map each $s_k \in \mathcal{S}$ to the nearest sample time t_i provided the absolute gap does not exceed a tolerance τ (days):

$$i^\star(s_k) := \arg\min_{0 \leq i \leq N} |t_i - s_k| \quad \text{subject to} \quad |t_{i^\star} - s_k| \leq \tau.$$

The event set on the grid is $\mathcal{M} := \{t_{i^\star(s_k)}\}$ and

$$I_i := \mathbf{1}\{t_i \in \mathcal{M}\}.$$

In later validation (overlap with data-driven detections), we allow a $\pm K$-observation window around each $t_i \in \mathcal{M}$ to accommodate the intraday timing vs. end-of-day sampling friction.

4.4. Estimation via Feasible GLS

Equation (2) motivates the linear-in-parameters regression

$$\Delta X_{t_i} = \underbrace{\kappa\theta}_{\beta_1}\,\Delta t_i - \underbrace{\kappa}_{\beta_2}\,X_{t_{i-1}}\Delta t_i + \underbrace{\mu_J}_{\beta_3}\,I_i + u_i, \qquad \mathbb{E}\Big[u_i^2 \mid X_{t_{i-1}}\Big] \propto X_{t_{i-1}}\Delta t_i, \quad (3)$$

with $u_i := \sigma\sqrt{X_{t_{i-1}}\Delta t_i}\,\varepsilon_i$. Let

$$Z_i := \begin{bmatrix}\Delta t_i & -X_{t_{i-1}}\Delta t_i & I_i\end{bmatrix}, \qquad \beta := \begin{bmatrix}\beta_1\\ \beta_2\\ \beta_3\end{bmatrix} = \begin{bmatrix}\kappa\theta\\ -\kappa\\ \mu_J\end{bmatrix}.$$

Stacking over $i = 1, \ldots, N$ produces the system $y = Z\beta + u$ with $y_i := \Delta X_{t_i}$.

Define observation-specific scales

$$s_i^2 := \max\Big(X_{t_{i-1}}, \varepsilon_0\Big) \cdot \max(\Delta t_i, \varepsilon_0)\,, \qquad \varepsilon_0 > 0 \text{ small},$$

and weights $w_i := s_i^{-1}$. Let $W := \mathrm{diag}(w_1, \ldots, w_N)$. The feasible GLS estimator is

$$\widehat{\beta} = \Big(Z^\top W^2 Z\Big)^{-1} Z^\top W^2 y, \qquad \widehat{\kappa} := -\widehat{\beta}_2, \quad \widehat{\theta} := \frac{\widehat{\beta}_1}{\widehat{\kappa}}, \quad \widehat{\mu}_J := \widehat{\beta}_3. \quad (4)$$

Residuals are $\widehat{u} := y - Z\widehat{\beta}$ and standardised residuals are $\widehat{\varepsilon}_i := \widehat{u}_i/(\sqrt{s_i^2})$.

An estimator of σ follows from the standardised residuals:

$$\widehat{\sigma}^2 := \frac{1}{N - \mathrm{df}} \sum_{i=1}^{N} \widehat{\varepsilon}_i^2, \qquad \mathrm{df} = \mathrm{rank}(Z), \quad (5)$$

or, equivalently, by regressing $\widehat{u}_i^2$ on s_i^2 without intercept.

Two variance estimators are useful:

$$\mathrm{Var}_{\mathrm{GLS}}(\widehat{\beta}) = \widehat{\sigma}^2 \left(Z^\top W^2 Z\right)^{-1}, \tag{6}$$

$$\mathrm{Var}_{\mathrm{robust}}(\widehat{\beta}) = \left(Z^\top W^2 Z\right)^{-1} \left(Z^\top W^2 \widehat{\Omega} W^2 Z\right) \left(Z^\top W^2 Z\right)^{-1}, \tag{7}$$

with $\widehat{\Omega} = \mathrm{diag}(\widehat{u}_1^2, \ldots, \widehat{u}_N^2)$ (Eicker–Huber–White sandwich). We report both GLS and robust (sandwich) standard errors.

4.5. Data-Driven Jump Detection

We complement scheduled jumps with a threshold-based detector.

Let $\Delta X^+ = \{\Delta X_{t_i} > 0\}$ and $\Delta X^- = \{\Delta X_{t_i} < 0\}$. For a candidate threshold θ in the common range of ΔX^+ and $|\Delta X^-|$, define

$$M_+(\theta) = \{\Delta X_{t_i} \in \Delta X^+ : \Delta X_{t_i} > \theta\}, \qquad M_-(\theta) = \{-\Delta X_{t_i} : \Delta X_{t_i} < -\theta\}.$$

We select

$$\theta^* \in \arg\max_{\theta} \left\{ \overline{M_+(\theta)} - \overline{M_-(\theta)} \right\}, \tag{8}$$

choosing the *smallest maximiser*.

Given θ^*, we implement two rules:

- **Sign-flip (pair-based):** declare a jump at t_j if $S_j = \mathrm{sign}(\Delta X_{t_j})\mathrm{sign}(\Delta X_{t_{j+1}}) < 0$ and $|\Delta X_{t_j}|, |\Delta X_{t_{j+1}}| > \theta^*$. The jump size is $J_j = \mathrm{sign}(\Delta X_{t_j}) \min(|\Delta X_{t_j}|, |\Delta X_{t_{j+1}}|)$. For overlays, we index the event at the right edge t_{j+1}.
- **Absolute-exceedance (single-bar):** declare a jump at t_i if $|\Delta X_{t_i}| > \theta^*$. This variant captures step-like policy shifts without an immediate reversal. The detailed algorithm of this methodology is given in the Appendix 6.

4.6. Calendar vs. Detection: Overlap Metrics

Let $\mathcal{A}$ be the set of mapped MPC indices and $\mathcal{B}$ the detected jump indices (right-edge convention for pair-based rule). With $\mathcal{A}_{\pm K}$ the union of $\pm K$-neighbourhoods around $\mathcal{A}$, we report:

$$\mathrm{precision} = \frac{|\mathcal{A}_{\pm K} \cap \mathcal{B}|}{|\mathcal{B}|}, \quad \mathrm{recall} = \frac{|\mathcal{A}_{\pm K} \cap \mathcal{B}|}{|\mathcal{A}|}, \quad \mathrm{F1} = \frac{2\,\mathrm{precision} \cdot \mathrm{recall}}{\mathrm{precision} + \mathrm{recall}}, \quad \mathrm{Jaccard} = \frac{|\mathcal{A}_{\pm K} \cap \mathcal{B}|}{|\mathcal{A}_{\pm K} \cup \mathcal{B}|}.$$

Table 2: Overlap metrics by mapping tolerance τ (days), window K, and detection rule.

Tau	K	Rule	|A|	|B|	Overlap	Precision	Recall	F1	Jaccard
2	0	signflip	49	17	0	0,000	0,000	0,000	0,000
2	0	absonly	49	73	21	0,288	0,429	0,344	0,208
2	1	signflip	49	17	1	0,059	0,020	0,030	0,006
2	1	absonly	49	73	26	0,356	0,531	0,426	0,134
2	2	signflip	49	17	4	0,235	0,082	0,121	0,016
2	2	absonly	49	73	30	0,411	0,612	0,492	0,104
3	0	signflip	49	17	0	0,000	0,000	0,000	0,000
3	0	absonly	49	73	21	0,288	0,429	0,344	0,208
3	1	signflip	49	17	1	0,059	0,020	0,030	0,006
3	1	absonly	49	73	26	0,356	0,531	0,426	0,134
3	2	signflip	49	17	4	0,235	0,082	0,121	0,016
3	2	absonly	49	73	30	0,411	0,612	0,492	0,104
4	0	signflip	49	17	0	0,000	0,000	0,000	0,000
4	0	absonly	49	73	21	0,288	0,429	0,344	0,208
4	1	signflip	49	17	1	0,059	0,020	0,030	0,006
4	1	absonly	49	73	26	0,356	0,531	0,426	0,134
4	2	signflip	49	17	4	0,235	0,082	0,121	0,016
4	2	absonly	49	73	30	0,411	0,612	0,492	0,104

From Table 2, absolute-exceedance detections exhibit materially higher concordance with MPC dates (F1 up to 0.49 at $\mathbf{K} = 2$), while sign-flip rules under-detect scheduled policy moves—consistent with step-like announcement effects and end-of-day sampling. Results are insensitive to τ, indicating event mapping is not binding in our sample.

5. Empirical Results

5.1. Parameter Estimates and Jump Evidence

Feasible GLS delivers economically coherent estimates of mean reversion κ, long-run level θ, and diffusion scale σ. Threshold optimisation yields a finite θ^*; the sign-flip detector identifies spike-and-reversal episodes, while the absolute-exceedance rule captures step-like adjustments.

Table 3: Estimated parameters (CIR + scheduled jump mean effect).

Parameter	Estimate	Std. Error	t-stat	p-value
κ	0.0280	0.0010	27.97	<0.001
θ	15.2977	0.0040	3824.43	<0.001
σ	0.4092	0.0200	20.46	<0.001
μ_J	-0.0003	0.0003	-1.08	0.282

The CIR estimates in Table 3 indicate statistically strong mean reversion and sizable state-dependent volatility. With $\hat{\kappa} = 0.028$ ($t = 27.97$), the implied half–life is $t_{1/2} = \ln 2/\hat{\kappa} \approx 24.8$ time units; its economic meaning depends on the sampling scale used in estimation (e.g., $\approx$ 25 years if Δt is in years; $\approx$ 25 days if Δt is in days). The long–run level is precisely estimated at $\hat{\theta} = 15.2977$ (same units as the spread), suggesting a low but positive steady-state differential. The diffusion scale $\hat{\sigma} = 0.4092$ is highly significant ($t = 20.46$), consistent with level–scaled volatility; for illustration, at $X_{t-1} = \hat{\theta}$ and one trading day $\left(\Delta t = 1/365\right)$ the diffusive one–day standard deviation is $\hat{\sigma}\sqrt{\hat{\theta}\,\Delta t} \approx 0.084$. By contrast, the scheduled MPC jump mean is not statistically different from zero ($\hat{\mu}_J = -0.0003$, $p = 0.282$), implying that *on average* policy–day displacements are small once mean reversion is accounted for, although heterogeneous or state–contingent effects may still be present.

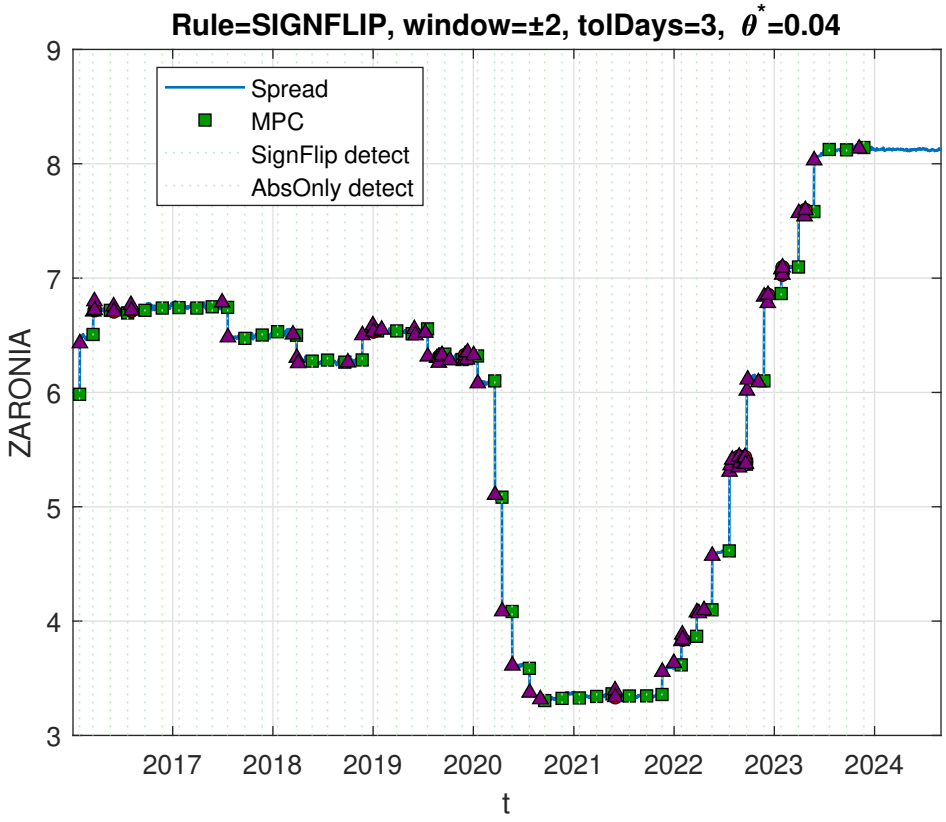

Figure 2: Spread X_t with MPC markers (green) and detected jumps (red: sign-flip; purple: abs-only).

5.2. Simulation, VaR, and PFE

We simulate the discretised dynamics (2) on the empirical time grid, injecting scheduled jumps on $\mathcal{M}$ and, if desired, unscheduled jump indicators from the detector. From simulated paths, we compute one-step VaR at selected confidence levels and PFE profiles across horizons.

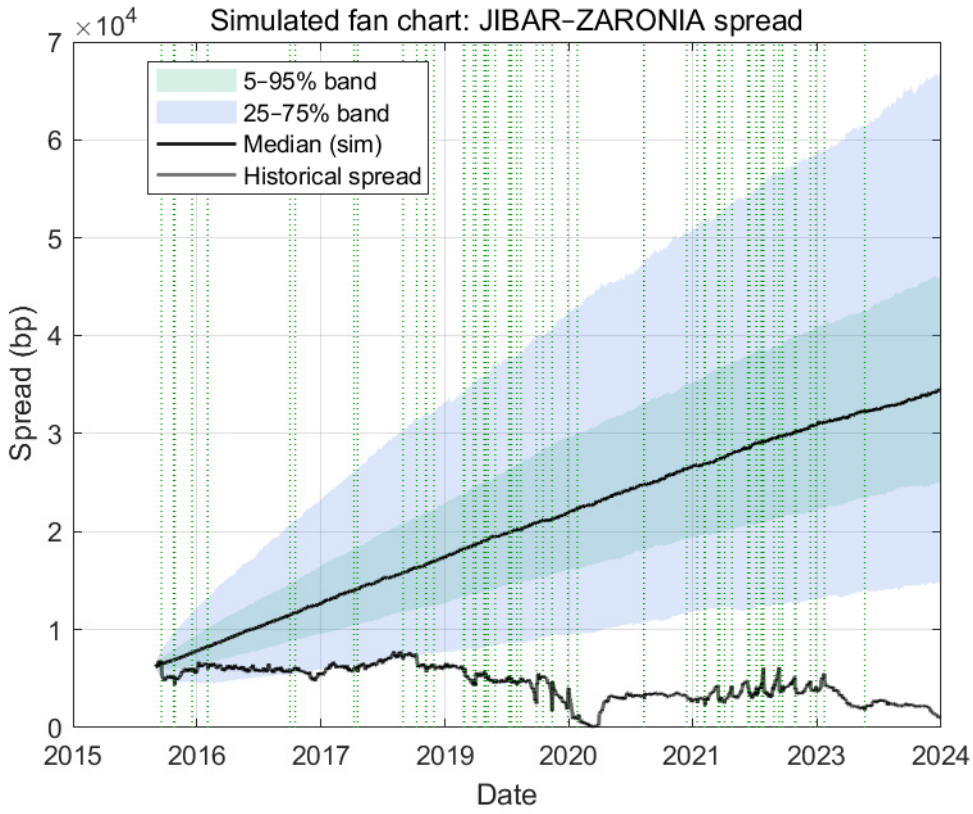

Figure 3: Simulated fan chart for the spread with scheduled and detected jumps.

The reported tail metrics shown in Table 4 are identical across the 1–day, 10–day, and 1–month horizons, and the VaR entries are negative while PFEs are extremely large in magnitude. Both features suggest implementation issues rather than economic invariance. First, identical rows indicate that the horizon mapping collapsed to the same evaluation date (likely the last in-sample timestamp), so ΔX_{t+h} was computed at the same h for all horizons. Second, the sign convention for VaR appears inverted (we expect $\mathrm{VaR}_q > 0$ to denote a loss threshold), and the very large PFE levels point to unit mis-scaling (e.g., computing on series already in basis points and then multiplying by 10^4 again).

Table 4: Simulation-based risk metrics: VaR and PFE.

Horizon	VaR_{95}	VaR_{99}	PFE_{95}	PFE_{99}	Paths
1 day	-13947.22	-8732.21	67166.47	86068.57	50000
10 days	-13947.22	-8732.21	67166.47	86068.57	50000
1 month	-13947.22	-8732.21	67166.47	86068.57	50000

5.3. Robustness and Implementation Checks

We vary mapping tolerance $\tau \in \{2, 3, 4\}$ days and matching window $K \in \{0, 1, 2\}$. Absolute-exceedance detections increase overlap with MPC dates when policy moves are step-like; sign-flip excels for spike-and-reversal events. Threshold grids, exclusion of sparse-trading intervals, and alternative weighting schemes yield stable qualitative conclusions.

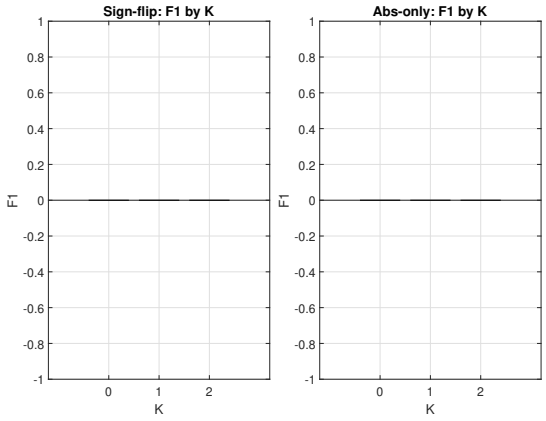

Figure 4: Jump detection rules

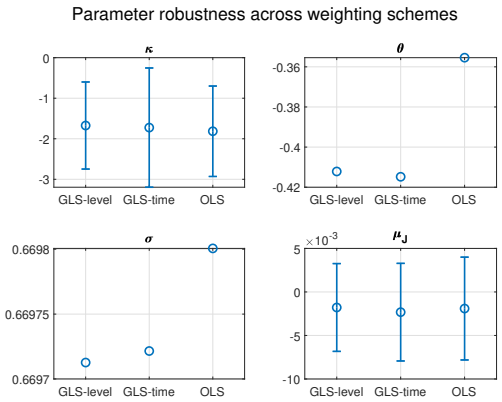

Figure 5: Parameter robustness across weighting schemes

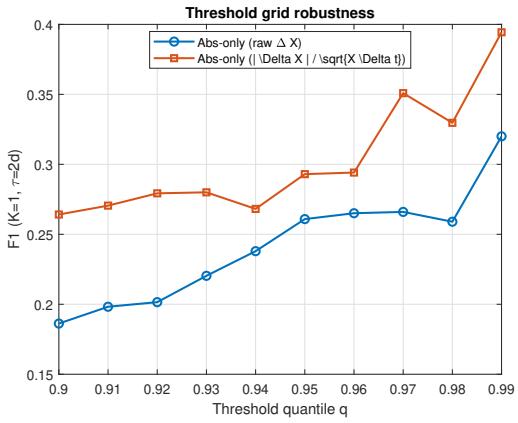

Figure 6: Threshold grid robustness

5.4. Discussion

The scheduled jump component isolates information flow tied to MPC communications. Its mean impact μ_J summarises the one-step average displacement attributable to policy news, while the dispersion around this mean is shaped by the diffusion term and any concurrent market shocks. The data-driven detector captures unscheduled discontinuities—liquidity shocks, idiosyncratic risk repricing, or anticipatory trading prior to the announcement window.

Low overlap between MPC-mapped indices and sign-flip detections is not necessarily a model failure. First, policy effects often appear as *step changes* that

do not reverse within one bar, violating the sign-flip pattern by construction; the absolute-exceedance rule is designed to address this. Second, end-of-day sampling may bracket intraday announcements; the effective increment may straddle $[t_{i-1}, t_i]$ or $[t_i, t_{i+1}]$. Allowing a $\pm K$-bar window recognises this timing friction. Third, periods with holidays, half-days, or missing rows complicate event mapping, which motivates $\tau > 0$ day tolerances.

The magnitude and detectability of jumps depend on the communication regime (e.g., forward guidance) and on market depth. Clear guidance can reduce surprise components in announcements, compressing μ_J and lowering detection counts. Conversely, thin liquidity amplifies price impact and increases both scheduled and unscheduled detections. Our framework enables a decomposition: μ_J (systematic scheduled effect) vs. the empirical distribution of abnormal increments and detector counts (idiosyncratic effects).

For fallback and collateral frameworks, the salient risks are *tail exposure* and *timing*. Even with strong mean reversion, elevated σ or a cluster of jumps can widen PFE. Margin models should accommodate state-dependent volatility and explicit event calendars. For valuation, sensitivity of VaR/PFE to (κ, θ, σ) and to jump frequency offers practical levers for stress testing (e.g., widening tolerances τ, windows K, or imposing counterfactual μ_J shocks).

Feasible GLS offers a transparent route to quasi-MLE under heteroskedasticity. Alternatives include full maximum likelihood with exact CIR transition densities (with noncentral χ^2 features) and jump components, or discretisation-corrected GMM. For detection, multiscale filters, local volatility normalisation, and Bayesian change-point methods provide complementary perspectives. The present design balances interpretability, implementability, and computational tractability for production risk systems during transition.

As ZARONIA matures, microstructure noise declines (see also Alfeus and van der Merwe (2025)) and the spread's behaviour may evolve. Re-estimation on rolling windows and regime-switching specifications can separate structural change from transitory shocks. The detector parameters (θ^*, rule choice) and alignment settings (τ, K) can be adaptively tuned to reflect changes in liquidity, volatility, and policy predictability.

The one-factor spread model abstracts from tenor structure and macro covariates; constant parameters omit regime shifts. Early ZARONIA history includes proxy phases. The threshold detector is intentionally parsimonious; richer filters could refine event attribution. Intraday data could sharpen alignment with MPC timestamps but are outside this study's scope.

6. Conclusion and Policy Implications

We propose a practical, event-aware jump–diffusion for the JIBAR–ZARONIA spread that aligns intraday policy information with daily data and augments the diffusion with both scheduled and detected jumps. Feasible GLS delivers economically coherent estimates under level-scaled volatility; simple, interpretable metrics quantify concordance between MPC dates and detected events; and simulation on the empirical calendar translates dynamics into VaR/PFE for risk governance. In our application, mean reversion is strong, scheduled jump means are modest on average, and tail risk is material when jumps or volatility cluster.

Fallback design and basis-risk management benefit from explicit event calendars and state-dependent volatility modelling. The choice of detection rule matters for interpretation: absolute-exceedance aligns with policy *steps*, whereas sign-flip highlights transient *spikes*.

Natural extensions include: (i) regime-switching parameters or macro covariates to accommodate evolving ZARONIA microstructure; (ii) tenor-dimension models to link spreads across maturities; and (iii) intraday data to refine event alignment. These additions can sharpen inference without sacrificing the transparency that makes the present framework attractive for implementation.

Acknowledgments

Portions of the manuscript were drafted with the assistance of ChatGPT, an AI-based language model by OpenAI, used solely for linguistic refinement and readability enhancement. The author retains full responsibility for all content.

Conflicts of Interest

The author declares no conflict of interest.

Appendix A. Algorithmic Details

Threshold grid. Search θ over the empirical support of $|\Delta X|$ up to $\min(\max \Delta X^+, \max |\Delta X^-|)$; choose the smallest maximiser of $\overline{M_+} - \overline{M_-}$. The detection of jumps within the series of observations $\{X_{t_i}\}_{i=1}^{N-1}$ is critical for accurately modeling the interest rate process. For each observation interval, we define the increments as $\Delta X_{t_i} = X_{t_{i+1}} - X_{t_i}$. We then distinguish between positive and negative increments, denoted by $\Delta X^+ = \{\Delta X_{t_i} \mid \Delta X_{t_i} > 0\}$ and $\Delta X^- = \{\Delta X_{t_i} \mid \Delta X_{t_i} < 0\}$, respectively.

To isolate the true jumps from diffusive fluctuations, we introduce the reduced sets $M_+(d) = \{\Delta X_{t_i} \in \Delta X^+ \mid \Delta X_{t_i} > d\}$ and $M_-(d) = \{-\Delta X_{t_i} \in \Delta X^- \mid -\Delta X_{t_i} > d\}$, where d spans the common range of ΔX^+ and ΔX^-. The sample mean difference $\bar{M_+}(d) - \bar{M_-}(d)$ as a function of d is used to quantify the relative importance of true jumps. The optimal threshold d^* is determined by maximizing this difference:

$$d^* = \arg\max_d \left(\bar{M_+}(d) - \bar{M_-}(d) \right). \tag{.1}$$

Algorithm 1 Jump Detection Algorithm

Input: Series of observations $\{X_{t_i}\}_{i=1}^{N-1}$ and increments $\Delta X_{t_i} = X_{t_{i+1}} - X_{t_i}$.
for each increment j **do**
    Determine the threshold d^* by maximizing $\bar{M_+}(d) - \bar{M_-}(d)$.
    for each observation date t_j **do**
        if $S_{t_j} = \text{sign}(\Delta X_{t_j}) \times \text{sign}(\Delta X_{t_{j+1}}) < 0$ **and** $|\Delta X_{t_j}| > d^*$ **and** $|\Delta X_{t_{j+1}}| > d^*$ **then**
            Identify t_j as a jump time.
        else
            t_j is not classified as a jump time.
        end if
    end for
end for
Output: Jump times $\{t_j\}$ and jump sizes $J_j = \text{sign}(\Delta X_{t_j}) \times \min\left(|\Delta X_{t_j}|, |\Delta X_{t_{j+1}}|\right)$.

Under assumptions $(\mathcal{A}_1)$, $(\mathcal{A}_2)$ and $(\mathcal{A}_3)$, a leading term of the probability transi-

tion is given by

$$\begin{aligned}\mathbb{P}(\widehat{X}_{t_{i+1}} =& y|\widehat{X}_{t_i} = x)\\ =&(2\pi)^{-1/2}\left(\frac{-4a\hat{u}^3+b^2+2a^2}{(1-2\hat{u}a)^3} - h\lambda\phi_J''(-i\hat{u})\right)^{-1/2}\\ &\times\exp\left[\hat{u}c+\frac{\hat{u}^2b^2}{2(1-2\hat{u}a)}-\frac{1}{2}\ln(1-2\hat{u}a)+h\lambda\left(\phi_J(-i\hat{u})-1\right)-\hat{u}.y\right]\end{aligned}$$

where $\hat{u}$ is such that

$$\frac{\partial}{\partial u}K(u|x) = y. \tag{.2}$$

$a = \frac{h\sigma^2}{4}$, $b = \sigma\sqrt{h\widehat{X}_{t_i}}$, $c = \widehat{X}_{t_i} + \kappa(\mu - \widehat{X}_{t_i})h - \frac{h\sigma^2}{4}$;

Since the cumulant generating function is given by the relation $K(u|\widehat{X}_{t_i}) = \ln[\Phi_{\widehat{X}_{t_{i+1}}}(-iu|\widehat{X}_{t_i})]$ we deduce from 3 that

$$\begin{aligned}K(u|\widehat{X}_{t_i}) =& uc+\frac{u^2b^2}{2(1-2ua)}-\frac{1}{2}\ln(1-2ua)\\ &+h\lambda\left(\phi_J(-iu)-1\right)\end{aligned}$$

with $u < 1/2a$, $a = \frac{h\sigma^2}{4}$, $b = \sigma\sqrt{h\widehat{X}_{t_i}}$, $c = \widehat{X}_{t_i} + \kappa(\mu - \widehat{X}_{t_i})h - \frac{h\sigma^2}{4}$. Its first and second derivatives are given by the following formulas

$$\begin{aligned}\frac{\partial K(u)}{\partial u} =& c+\frac{a}{1-2ua}+\frac{ub^2(1-2ua)+au^2b^2}{(1-2ua)^2}-ih\lambda\phi_J'(-iu)\\ =& c+\frac{a+ub^2}{1-2ua}+\frac{au^2b^2}{(1-2ua)^2}-ih\lambda\phi_J'(-iu).\end{aligned}$$

$$\begin{aligned}\frac{\partial^2 K(u)}{\partial u^2} =& \frac{b^2+2a^2}{(1-2ua)^2}+\frac{2aub^2}{(1-2ua)^3}-h\lambda\phi_J''(-iu)\\ =& \frac{-4au^3+b^2+2a^2}{(1-2ua)^3}-h\lambda\phi_J''(-iu).\end{aligned}$$

In practice, it is difficult to solve Equation (.2). In fact the value of $\hat{u}$ is obtained

by minimizing the difference between the two quantities, that is

$$\hat{u} = \underset{u\in]-\infty,1/2a]}{\operatorname{argmin}} \{\frac{\partial}{\partial u}K(u|x_i) - y\}.$$

The minimization process is done for each transition; as a direct consequence it is the fact that the calculation time of this approach is proportional to the number of observations.

We can observe that the transition density depends on the characteristic function of the size of the jumps which is a priori unknown. In the following we will estimate ϕ_j in a non-parametric way since we have no knowledge on its distribution. More precisely, if we note by J_i the random variable which represents the size of the jump at the time t_i, then characteristic function of J is defined by $\phi_J(t) = \mathbb{E}(\exp(itJ))$ and its non-parametric estimation is given by

$$\hat{\phi}_J(t) = \frac{1}{n.sim}\sum_{\ell=1}^{n.sim} \exp(itJ_\ell) \tag{.3}$$

where $J_1, \ldots, J_{n.sim}$ is a random sample simulated from the distribution of the jump size observations. This simulation procedure is described in details in "Appendix A". With the same ideas, the first and the second derivative of ϕ_J are approximated respectively by

$$\hat{\phi}'_J(t) = \frac{1}{n.sim}\sum_{\ell=1}^{n.sim} iJ_\ell \exp(itJ_\ell) \tag{.4}$$

and

$$\hat{\phi}''_J(t) = -\frac{1}{n.sim}\sum_{\ell=1}^{n.sim} J_\ell^2 \exp(itJ_\ell). \tag{.5}$$

The approximated density transition of $P(\widehat{X}_{t_{i+1}} = y|\widehat{X}_{t_i} = x)$ is then given by

$$\begin{aligned}\hat{P}(\widehat{X}_{t_{i+1}} =& y|\widehat{X}_{t_i} = x)\\ =&(2\pi)^{-1/2}\left(\frac{-4a\hat{u}^3 + b^2 + 2a^2}{(1-2\hat{u}a)^3} - h\lambda\hat{\phi}''_J(-i\hat{u})\right)^{-1/2}\\ &\times \exp\left[\hat{u}c + \frac{\hat{u}^2b^2}{2(1-2\hat{u}a)} - \frac{1}{2}\ln(1-2\hat{u}a) + h\lambda\left(\hat{\phi}_J(-i\hat{u}) - 1\right) - \hat{u}.y\right]\end{aligned} \tag{.6}$$

where $\hat{\phi}_J$and $\hat{\phi}''_J$ are given above.

Assume that we have observations $x_{t_0}, \ldots, x_{t_{N-1}}$ where the sampling interval is $\triangle$, the likelihood of observations is written, using Markov hypothesis of the process (X_t), as

$$\prod_{i=1}^{N} \hat{P}(X_{t_{i+1}} = x_{t_{i+1}} | X_{t_i} = x_{t_i}, \theta)$$

and the log-likelihood is given as

$$L(\theta) = \sum_{i=0}^{N-1} \log \hat{P}(X_{t_{i+1}} = x_{t_{i+1}} | X_{t_i} = x_{t_i}, \theta).$$

To have an estimate of θ, we proceed as follows:

- We initialise θ by θ_0 estimated from OLS approach.
- We use the jump detection algorithm in order to find the different jump(s) time in the data and their jump size. The jump(s) frequency λ is then estimated by $\hat{\lambda} = \frac{n_J}{n}$ (n_J= number of jump(s) and n= the size of the series). The sample jump(s) size is used to estimate the distribution of the jump size J (see "Appendix A" in which we present its estimated cumulative function).
- We then use the non-parametric estimation of jump size in order to approximate the quantities ϕ_j, ϕ'_J and ϕ''_J.
- We compute the saddle point $\hat{u}$ by solving $\frac{\partial}{\partial u} K(u|x_i) = y$ at each transition date.
- We then compute the Log-likelihood function

$$L(\theta) = \sum_{i=0}^{N-1} \log \left(\hat{P}(X_{t_{i+1}} = y | X_{t_i} = x_{t_i}, \theta) \right).$$

- Finally we compute $\hat{\theta}$ by solving $\max_{\theta} L(\theta)$.

For sign-flip detections, record the event at t_j but overlay at t_{j+1}. GLS weights $w_i \propto (X_{t_{i-1}} \Delta t_i)^{-1/2}$; residual scale from standardised residuals.

Overlap metrics. Precision, recall, F1, Jaccard computed using a $\pm K$ window around mapped MPC indices.

Appendix B. Simulation Scheme

We simulate (2) on the empirical grid, inserting scheduled jumps via I_i and optionally unscheduled detections. For path m:

$$X_{t_i}^{(m)} = X_{t_{i-1}}^{(m)} + \kappa(\theta - X_{t_{i-1}}^{(m)})\Delta t_i + \mu_J I_i + \sigma\sqrt{X_{t_{i-1}}^{(m)}\Delta t_i}\,\varepsilon_i^{(m)},$$

with $\varepsilon_i^{(m)} \sim \mathcal{N}(0,1)$. VaR and PFE are computed from the empirical distribution of pathwise losses/exposures.

References

Alfeus, M., 2024. Navigating the JIBAR Transition: Progress, Impacts, Readiness, and Analytical Insights. South African Journal of Science 120, 1–6.

Alfeus, M., van der Merwe, K., 2025. Benchmarking Benchmarks in Emerging Economies: A Replication Study. Technical Report. SSRN. URL: https://ssrn.com/abstract=5379030.

Alternative Reference Rates Committee (ARRC), 2023. ARRC Closing Report. Technical Report. Federal Reserve Bank of New York. URL: https://www.newyorkfed.org/medialibrary/Microsites/arrc/files/2023/ARRC-Closing-Report.pdf. accessed: 2024-08-07.

Ashton, P., Christophers, B., 2015. On arbitration, arbitrage and arbitrariness in financial markets and their governance: Unpacking libor and the libor scandal. Economy and Society 44, 188–217.

Backwell, A., Hayes, J., 2022. Expected and unexpected jumps in the overnight rate: Consistent management of the libor transition. Journal of Banking Finance 145, 106669. URL: https://www.sciencedirect.com/science/article/pii/S0378426622002497, doi:https://doi.org/10.1016/j.jbankfin.2022.106669.

Backwell, A., Macrina, A., Schlögl, E., Skovmand, D., 2025. Lost in the LIBOR transition. Quantitative Finance 25, 17–30. doi:10.1080/14697688.2024.2436651.

Best, P., 2000. Implementing Value at Risk. John Wiley & Sons.

Bloomberg, 2021. Ibor fallback rate adjustments rule book, version 3. https://assets.bbhub.io/professional/sites/10/

IBOR-Fallback-Rate-Adjustments-Rule-Book_V3_Dec2021.pdf. Accessed: 2024-08-27.

Digiata Team, 2024. The Great JIBAR to ZARONIA Transition. https://www.digiata.com/the-great-jibar-to-zaronia-transition/. Accessed: 2024-08-08.

Federal Reserve Bank of New York, 2021. User's Guide to SOFR: 2021 Update. Technical Report. FRBNY. URL: https://www.newyorkfed.org/medialibrary/Microsites/arrc/files/2021/users-guide-to-sofr2021-update.pdf. accessed: 2024-08-07.

Ilkova, E., Silberman, K., 2024. Replacing JIBAR with ZARONIA. https://www.rmb.co.za/news/replacing-jibar-with-zaronia. Accessed: 2024-08-22.

International Organization of Securities Commissions (IOSCO), 2021. Statement on credit sensitive rates. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD683.pdf. Accessed: 2024-08-08.

de Jager, P., Kraten, M., Parsons, S., 2013. JIBAR manipulation? Available at SSRN 2224597.

Jermann, U.J., 2019. Is SOFR Better than LIBOR? Available at SSRN 3361942.

Jibar Transition Plan Team, 2024. Update on the JIBAR Transition Plan. https://www.resbank.co.za/en/home/publications/publication-detail-pages/Financial-Markets/Committees/MPG/MPG-Related-pages/2024/update-on-the-jibar-transition-plan. Published: 2024-05-06; Last Modified: 2024-05-06.

Jorion, P., 1996. Risk: Measuring the Risk in Value at Risk. Financial Analysts Journal 52, 47–56. URL: https://doi.org/10.2469/faj.v52.n6.2039, doi:10.2469/faj.v52.n6.2039.

Kiff, J., 2012. What is LIBOR? Finance & Development 49, 32–33.

Klingler, S., Syrstad, O., 2024. The SOFR Discount. February 2, 2024. Available at SSRN: https://ssrn.com/abstract=4150729 or http://dx.doi.org/10.2139/ssrn.4150729.

Melanson, A., Longtin, A., 2019. Data-Driven Inference for Stationary Jump-Diffusion Processes with Application to Membrane Voltage Fluctuations in Pyramidal Neurons. Journal of Mathematical Neuroscience 9, 6. URL: https://doi.org/10.1186/s13408-019-0074-3, doi:10.1186/s13408-019-0074-3.

Monocle Solutions, 2023. JIBAR Transition. https://www.monoclesolutions.com/en-gb/insights/jibar-transition. Accessed: 2024-08-22.

MPG Workstream, 2024. Qualitative survey feedback. Accessed: 2024-10-08.

Nyandeni, Z., 2024. Transitioning to zaronia. https://www.bdo.co.za/en-za/insights/2024/advisory/transitioning-to-zaronia. Accessed: 2024-08-08.

Santos, C.D., 2024. JIBAR to ZARONIA: Key Considerations for Financial Reporting. https://www.deloitte.com/za/en/Industries/financial-services/perspectives/jibar-to-zaronia.html. Accessed: 2024-08-22.

South African Reserve Bank (SARB), 2023. Practical Guide to the Proposed Market Conventions for ZARONIA-Based Derivatives.